LBNL-6104E

Ernest Orlando Lawrence Berkeley National Laboratory

Smart Grid as a Driver for Energy-Intensive Industries: A Data Center Case Study

Venkata Ganti and Girish Ghatikar

Lawrence Berkeley National Laboratory

December 2012

Presented at
Grid-Interop 2012,
Irving, TX,
December 3-6, 2012,
and published in
the Proceedings

Smart Grid as a Driver for Energy-Intensive Industries: A Data Center Case Study

Venkata Ganti
VSGanti@lbl.gov
Lawrence Berkeley National Laboratory
1 Cyclotron Road, 90R1121, Berkeley, CA 94720

Girish Ghatikar
GGhatikar@lbl.gov
Lawrence Berkeley National Laboratory
1 Cyclotron Road, 90R1121, Berkeley, CA 94720

Keywords: Data Centers, Smart Grid, Load Migration, Demand Response, Enabling Technologies, Grid Integration

Abstract

The Smart Grid facilitates integration of supply- and demand-side services, allowing the end-use loads to be dynamic and respond to changes in electricity generation or meet localized grid needs. Expanding from previous work, this paper summarizes the results from field tests conducted to identify demand response opportunities in energy-intensive industrial facilities such as data centers. There is a significant opportunity for energy and peak-demand reduction in data centers as hardware and software technologies, sensing, and control methods can be closely integrated with the electric grid by means of demand response.

The paper provides field test results by examining distributed and networked data center characteristics, end-use loads and control systems, and recommends opportunities and challenges for grid integration. The focus is on distributed data centers and how loads can be "migrated" geographically in response to changing grid supply (increase/decrease). In addition, it examines the enabling technologies and demand-response strategies of high performance computing data centers. The findings showed that the studied data centers provided average load shed of up to 10% with short response times and no operational impact. For commercial program participation, the load-shed strategies must be tightly integrated with data center automation tools to make them less resource-intensive.

1. INTRODUCTION

The recent boom in cloud computing has led to the burgeoning of distributed and co-location data centers. This growth has led to the development and deployment of high-density storage systems and high performance computing (HPC) systems that tend to be more energy efficient but consume more energy than the legacy systems [1]. Studies have estimated that computing efficiency, measured in computations per kilowatt-hour, is doubling every 1.5 years [2]. Despite these efficiency improvements in cooling and IT equipment, the net energy consumption of data centers is still on the rise. According to a 2007 U.S. Environmental Protection Agency (EPA) report, 20% of the U.S. data center energy use is in the Pacific region alone [3][4].

The maturity in cloud computing and virtualization technologies has enabled data centers to dynamically migrate load to distributed (also called *co-location*) data centers for disaster recovery and backups. With slight modifications, such infrastructure can be integrated with the Smart Grid. Dynamically migrating loads and lowering the demand on peak days or when the electricity costs are much higher can improve grid reliability and reduce data center operational costs. Such unique features within the norms of the data center Service Level Agreements (SLAs) improve grid reliability. With acceleration and investment in the U.S. Smart Grid deployment, U.S. data centers can benefit from the value provided by the Smart Grid and participate in Demand Response (DR) programs. This paper expands on field tests of these concepts in previous studies [5][6].

The IT equipment loads of a smart grid-integrated or simply "grid integrated", data center can be migrated to a location where the energy is cheaper or more readily available. Load shed and shift strategies have been widely used for commercial and industrial DR [8][9]. Data centers present a unique opportunity to shift or migrate load from one location to another, be it a different grid, utility, or state, or even a different country. If a utility has a higher percentage of variable generation resources such as wind, there are times when there is excess generation due to higher wind speeds. Building and maintaining grid-level energy storage systems is expensive and complex. A rational alternative is to use energy at locations with excess renewable generation by migrating computational jobs from a geographically distant data center to the other. Considering these advantages, load migration strategy with the required automation can be an excellent strategy for distributed and networked data centers.

1.1. Background

The Demand Response Research Center (DRRC) at the Lawrence Berkeley National Laboratory (LBNL) has conducted research to identify the potential for data centers to participate in demand response. In 2010, DRRC published a report that characterized data centers based on operational characteristics, equipment and end-uses, energy usage, and load profiles [7]. The findings highlighted some of the key characteristics of data centers, such as the synergy between IT and site. It was concluded that largest potential for load reduction can come from automating DR strategies in IT equipment, such as server load management and load migration in non-mission-critical data centers [5].

As a follow-up to this study, in 2012, DRRC conducted a series of field tests of DR strategies at four data centers: (1) NetApp, a backup storage data center at their Sunnyvale, California, campus, (2) Lawrence Berkeley National Laboratory (LBNL), Building 50 data center, (3) the University of California (UC) Berkeley Data Center, and (4) San Diego Super Computer Center. In August 2012, LBNL published a report [6] on the results from the field tests for various DR strategies. In this paper, we take a closer look at the load migration strategy, which we believe presents the largest opportunity for load shed during demand response. The term *load migration* refers to geographic shifting of the computing load from one grid-responsive data center to another.

2. DATA CENTERS AND GRID INTEGRATION

When data centers become integrated with the Smart Grid, they are not only "self-aware" to meet local needs, they also become "grid-aware" to respond to changing grid conditions (e.g., price or reliability) and gain additional benefits resulting from incentives and credits and/or lowered electricity prices and other markets. Data centers can be integrated with the grid by programming and the enabling technologies and data center automation software to listen to real-time energy pricing information from wholesale or retail energy markets. To identify areas of grid integration, the data center has to be characterized to understand how it fits within similar data centers.

3. DATA CENTER CHARACTERIZATION

Data centers are energy-intensive industrial facilities that house a collection of IT equipment servers and storage and network devices in a dedicated space [5][6]. The IT equipment is supported by power, cooling, and lighting systems, which are referred to as the *site infrastructure.* The data center characterization involves collecting and analyzing technical and operational information. Understanding the key operational characteristics of a data center is essential for developing a preliminary list of DR opportunities, along with their feasibility and potential for automation. This characterization task consists of four key attributes: data center functions, enabling technologies, load profiles, and computational job characterization.

1. ***Data Center Functions*** refers to the type of service offered by a data center**.** For example, web servers, file or media storage servers, database systems, and high performance computing are some of the most common services offered by modern day data centers.
2. ***Enabling Technologies*** specific to data centers can provide information to the data centers to facilitate DR program participation. These technologies provide real-time management and control of IT equipment, cooling, and monitoring of temperature and humidity conditions for cooling and air management. The technologies that manage computing loads also provide data to better characterize the field test results.
3. ***Load Profile*** refers to the categorization of various end-use loads based on its function. End-use loads in a data center environment can be categorized into three broad categories: (1) IT equipment, (2) cooling or site infrastructure, and (3) support loads, which consist of an uninterruptible power supply (UPS) and lighting.
4. ***Computational job characterization*** involves characterizing the jobs performed by the data center servers to develop a deep understanding of the type of services being run at the center. This process identifies potential jobs to migrate to a co-location data center during a DR event. Some jobs require local resources and may have different computational needs than those available at a co-location data center. In such special circumstances, computational jobs cannot be migrated. However there may be a chance to reschedule such jobs to run before or after the DR event.

The approach used to characterize data centers has been well described in the LBNL Phase 2 report. Data centers similar to the ones that were part of this study can apply this approach to identify the applicable DR strategies.

4. DEMAND RESPONSE STRATEGIES

Data center DR opportunities depend on several factors, including the institutional and technical capabilities identified in the previous section. Data centers with virtualization and enabling technologies for servers, storage, and networking equipment have the largest potential for DR. Similarly, by automating the cooling system response to IT equipment provides an opportunity for synergistic load reduction. Raising temperature set points and lighting strategies can be the first-order DR strategies, which have been well studied [5][6].

Data center managers may perceive that some strategies are applicable for energy efficiency; however, raising the bar and temporarily reducing service levels without impact to operations can achieve further incremental benefits. These DR strategies generally fall into the categories of *load shedding* (dropping load completely) and *load-shifting* (moving load from peak to off-peak periods). These strategies can apply to both IT equipment and site infrastructure.

The specific data center DR strategies vary, depending on their services and the type of IT equipment. For example, in a data storage system, the potential DR opportunities include rescheduling of storage jobs such as backup processes and idling or powering off filer heads and their associated storage shelves. While a stand-alone data center can use many DR strategies [5][6], this paper emphasizes the integration of a decentralized data center network.

Distributed and networked data centers provide a unique ability to continue normal operations by migrating the loads as grid or price conditions change. Such load migration strategies offer significant promise for data centers to benefit from the existing disaster recovery infrastructure and gain additional value by improving grid reliability and lowered operational expenses.

5. LOAD MIGRATION STRATEGIES

Most cloud-based data centers maintain fully networked and distributed locations on different electrical grids or geographic locations as back up for disaster recovery. Depending on the service level agreements (SLAs) and uptime requirements, each data center site in the cloud maintains its own backup generation or energy storage system to come into action in the event of a power grid failure. LBNL discussed with data center experts the emerging technologies currently available or in development that could allow temporary load migration of data center IT equipment loads outside a region that is experiencing a DR event. Because of this shift, IT equipment could be shut down or enabled for intelligent power management. Although this is primarily an IT infrastructure strategy, the shift in IT loads would reduce supporting site infrastructure (cooling) loads as well. Data centers that participate in DR using this strategy would likely need advance notice of the need for load migration for planning and coordination. The target loads for this DR strategy are computing nodes, processors, hard drives, networking components, and internal system thermal management systems

5.1. Concept of Capacity Reservation in HPC systems

A *capacity reservation block* is a policy that is created in the job scheduler to block a certain number or percentage of nodes from running IT jobs. When it comes into effect, IT jobs running for that percent of the nodes are halted or killed, and new jobs are prevented from starting on those nodes. After the reservation block is removed, the percent of nodes can either be sent to idle mode or shut down completely. Figure 1 shows an example of the capacity reservation concept.

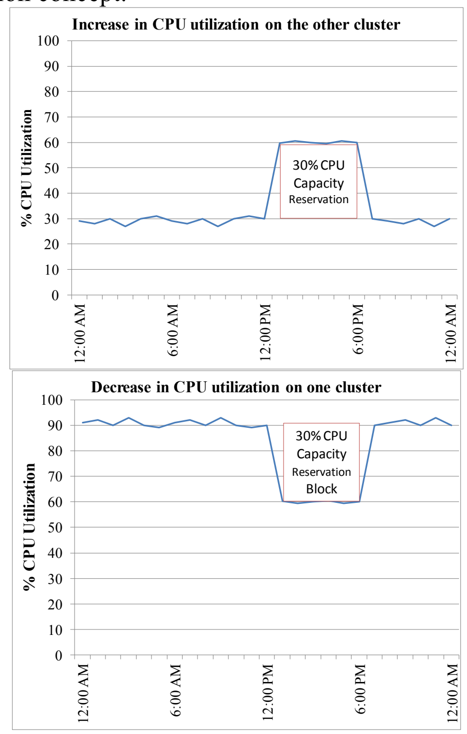

Figure 1. Example of a Reservation Block Allocation

The considerations for deploying this strategy vary, depending on the availability of resources at the data center accepting the migrated computational jobs. If the participating IT equipment at both the co-location data centers vary by type/specifications, they are called *heterogeneous systems.* If the IT equipment is completely mirrored at both data centers, then they are called *homogeneous systems.* The load migration capabilities of homogenous and heterogeneous systems differ and are discussed in detail in the following sections.

5.2. Sequence of Operations for Load Reduction

A DR strategy must be implemented to achieve the desired load shed from any system without affecting the services and/or life of the equipment. There may be differences in the enabling technologies used and their capabilities, but the overall procedural algorithm remains the same as described herein. The sequence of operations and related information is presented in more detail in the LBNL publication [6]. These general guidelines are applicable for both homogeneous and heterogeneous HPC clusters.

As a first step in the process, the status of CPU utilization and power consumption on both ends of the HPC cluster should be monitored. Depending on the availability of resources and jobs in the job scheduler queue, a percentage of capacity reservation should be selected and placed on the HPC system housed in the data center that will participate in the DR event. This capacity reservation will prevent new jobs from starting during the DR event.

Once the DR event starts, the capacity reservation will come into effect and push the applicable percentage of HPC nodes into idle mode, resulting in a load shed. If the infrastructure can turn down, there will be a magnifying of load shed. Depending on the criticality of jobs currently running, they can either be killed or allowed to decay over time, resulting in a slower load-shed response. When possible, a percentage of the computing nodes can be powered down completely to achieve higher load shed, keeping in mind the standard operating procedures of a participating data center. The capacity reservation will stay in effect until the end of the DR event period.

5.3. Sequence of Operations for Recovery

Once the DR event ends, the HPC clusters need to be restored to the normal state of operation without any impact to the equipment. First, the capacity reservation on the job scheduler should be removed, and the idling or powered down nodes should be restarted.

If all the systems are restarted at the same time, the nodes that are trying to initiate the booting process will create a new peak. To prevent the occurrence of spikes in peak demand, compute nodes should be turned on sequentially by inducing a timed delay between each other. This staging or staggering of the booting process ensures a reliable restoration without straining or blocking the network tunnels. Once all the nodes are up and running, the job scheduler adds these nodes to the pool of available resources, and it can start accepting new compute jobs.

5.4. Field Test Results

LBNL conducted field tests for the load migration strategy at three data center sites: the LBNL Building 50 data center, the UC Berkeley data center, and the San Diego Super Computer center (SDSC). The tests included migrating computation jobs on the Shared Research Computing Services Cluster (ShaRCS) system, which is a homogeneous cluster co-hosted by UC Berkeley and SDSC. The computing cluster located at the UC Berkeley data center is called *Mako,* while its sister cluster at SDSC is called *Thresher.* Table 1 summarizes the results.

The *response period* refers to the time taken for the cluster to shed the desired percentage of load. The table shows the time taken by the clusters to respond and shed 5% and 10% of their load. The table also shows the recovery period, which is the amount of time taken by clusters to return to normal operational modes. The first two tests show the results when the jobs are migrated fully before idling or shutting down the nodes. In case of load migration in heterogeneous clusters, jobs are allowed to decay slowly over a period of time. This decay process is the reason behind the longer response times, as shown in Table 1.

Table 1: Summary of Test Results

DR Strategy	Response Period (min)		Recovery Period (min)
	5% shed	10% shed	
Load migration in Homogenous Cluster (Idling)	2	6	2
Load migration in Homogenous Cluster (Shutdown)	3	7	10
Load migration in Heterogeneous Cluster (Decay)	147	175	15

To meet the industry testing standards, the clusters were loaded with High Performance Linpack (HPL) benchmarking jobs to raise their CPU utilization level to 95% and 40% on Mako and Thresher, respectively. A capacity reservation block of 30% was placed on the Mako cluster, and the corresponding HPL jobs were migrated to run on Thresher. Figure 2 shows the linear correlation between CPU utilization and its corresponding power draw.

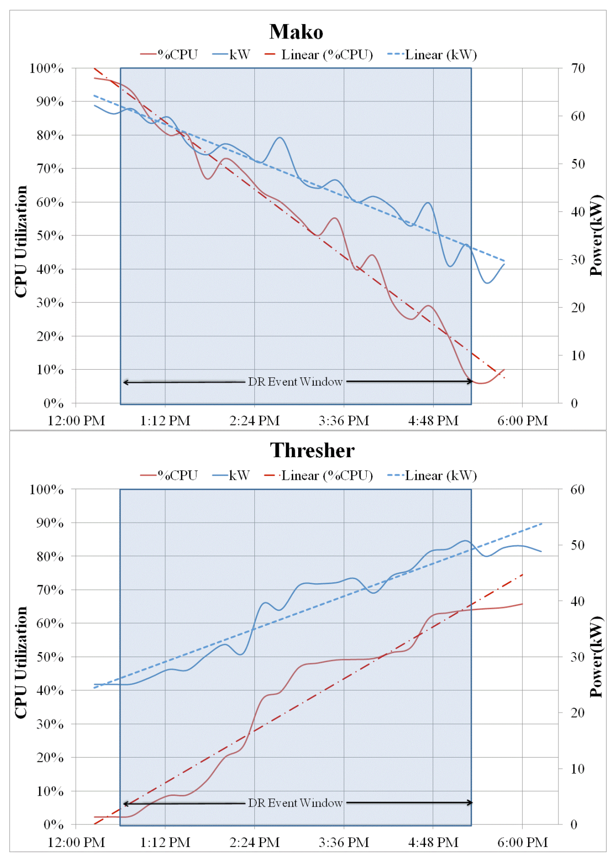

Figure 2: ShaRCS Load Migration Test Results

Once the jobs were migrated, the Mako nodes, which were not processing any jobs, were pushed into idle mode, resulting in the reduction of 8.7 kilowatts (kW), which corresponds to 14% of the entire Mako system demand. This achievement was possible in an impressive four minutes, making this a groundbreaking strategy for fast-DR, ancillary services, and current DR program offerings.

Summary of Homogeneous Load Migration Strategy

The load migration DR strategy can be categorized for IT equipment such as servers, storage, and networking systems. When the DR event is initiated, the following steps must be followed in the same sequence to participate in the event without causing any operational or hardware issues.

1. Set an appropriate capacity reservation (%) on Cluster 1 to prevent starting of new jobs during the DR event.
2. At the start of the DR event, the capacity reservation is in effect, and it pushes the HPC nodes into idle mode.
3. Shutdown the IT equipment to achieve higher load shed if it meets the standard operating procedures of a participating data center.
4. Hold Cluster 1 in this state until the end of the event.

After the DR event is completed, the data center operations can be restored to normalcy by following the sequence of recovery operations presented below.

1. After the DR event, lift the capacity reservation to return the system to its normal operational state.
2. Put the idled nodes back in active mode to accept new jobs.
3. Restart the shutdown nodes and put them in active mode to accept new jobs.

As a best practice, a time delay must be given between the start-up for each node. In the field tests conducted by the LBNL team, the computational nodes were stateless and required booting from the network. Systems which do not store the operating system on a local hard disk but boot from the network are called stateless systems. A four-second delay was induced to stage the booting of the nodes without straining the network or causing a spike in the power draw due to starting them all simultaneously. The load migration strategy is well studied and documented [6].

A similar test was performed on the Lawrencium cluster, which is co-hosted by LBNL Building 50 data center and SDSC. The Lawrencium cluster is a heterogeneous computing cluster consisting of LR-1 at SDSC and LR-2 at the LBNL data centers. In this case, the demand shed of 9 kW (17% of the LR-1 cluster) was achieved. There was an increase in the power consumption of LR-2 at the same time due to the increase in CPU utilization resulting from the jobs that were migrated. In-depth analysis and results of both these tests were presented in the LBNL data center field study report [5]. This report also related the findings to the common DR framework specified by the National Institute of Standards and Technology (NIST) and Organization for the Advancement of Structured Information Standards (OASIS) committees [10][11][12].

Summary of Heterogeneous Load Migration Strategy

The sequence of operations for heterogeneous load migration is very similar to homogeneous load migration; however, some additional factors must be considered.

Prior to the placement of a load capacity reservation block on the clusters, the CPU utilization and power consumption on both the clusters have to be measured in real time. This enables data center operators to determine if migrating the load to a different location will be cost effective and not cause operational issues.

In the test conducted by LBNL, jobs currently running were allowed to finish processing and decay over a three-day period. After all the jobs are drained, the LR-1 nodes run in idle mode until the end of the DR event. Nodes can be shut down to achieve higher load shed, provided the strategy meets the data center's standard operating procedures.

The sequence of operations for recovery is the same as that for the homogeneous system. Figure 3 shows the results from the heterogeneous load migration tests conducted on the Lawrencium cluster. It can be seen that the load on LR-1 decays slowly over a period of few hours. The Load on LR-22 increased slightly because of the new jobs which were migrated from LR-1.

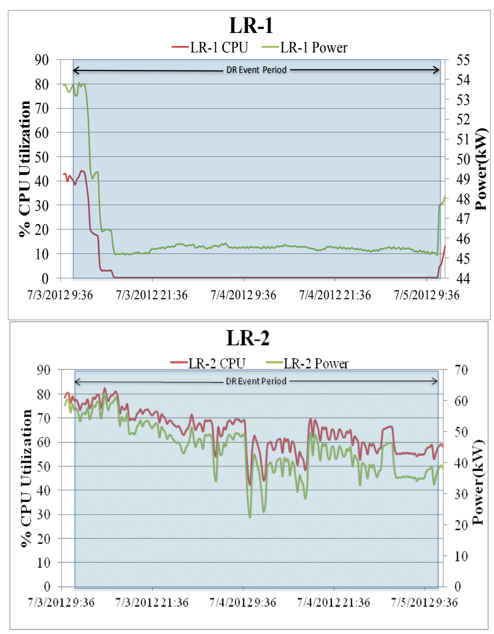

Figure 3: Lawrencium Load Migration Test Results

5.4.1. CPU Utilization and Power Correlation

Both the Shared Research Computing Services Cluster (ShaRCS) and Lawrencium (LR) Cluster load migration tests show a strong correlation between CPU utilization and power consumption. When the computing nodes are performing complex computations involving many parallel processing nodes and memory input/output operations, the CPU utilization generally increases, resulting in a larger power draw. This has been theoretically reported in many publications, but this is the first time that those claims have been validated with actual field tests and analysis of results.

6. ENABLING TECHNOLOGIES FOR GRID INTEGRATION

To facilitate the communication between the data center and the local electricity service provider, data center enabling technologies require the development of software to integrate data center operations with the electric grid. This requires a core understanding of IT resource management applications and their control capabilities and algorithms.

Automation of Enabling Technologies

Implementing DR strategies at data centers can be a tedious task, and it requires the special attention of data center IT and facilities managers. Several distributed energy management and control systems (EMCSs) are currently on the market and are primarily used for monitoring and implementing energy-efficiency measures. These systems regulate operation of HVAC, lighting, and related electrical systems in an integrated fashion. Communication building control protocols such as BACnet®, Modbus®, and LonTalk® allow EMCS to communicate with site infrastructure equipment. These protocols are important to understand and can be programmed to communicate any efficiency or potential DR strategy, as well as oversee technology interoperability within data centers. In many cases, such EMCS or supervisory control and data acquisition (SCADA) systems can be preprogrammed to manage data center support loads in response to a DR event notification.

OpenADR Integration with Control Systems

The first step toward automating demand response strategies is to use open architectures for integrating data center energy management and control systems with the information from the power grid. National Smart Grid standards such as OpenADR can be used to receive grid reliability and pricing signals from local utilities or the grid operators (ISOs). Using Transmission Control Protocol over Internet Protocol (TCP/IP) and eXtensible Markup Language (XML) would enable open, standards-based information exchange within a data center's virtualization network and interoperability (as well as integration) with the Smart Grid. Technologies that integrate site and IT infrastructure would be useful to provide a single source of information for integrated implementation of DR strategies.

7. LINKS TO THE GWAC INTEROPERIBILITY FRAMEWORK

The work presented in this paper meets the organizational and informational categories of GridWise Architecture Council's (GWAC) interoperability context-setting framework. Data centers can participate in DR programs to meet their business operation goals and set policies based on internal information to optimize the performance and utilization of a distributed and networked data center. From a technological perspective, integrating data centers with the grid can meet the interoperability and other technical aspects of the GWAC's interoperability framework [12].

8. CONCLUSIONS AND NEXT STEPS

The study results show that IT loads can be turned off manually in a DR event in less than eight minutes and achieve an average load shed of 10%, which is noticeable at the whole building level. This makes data centers excellent candidates to participate in Auto-DR programs and integrate with OpenADR for retail and wholesale DR markets.

Future research direction must look at the potential of load migration strategies through the development of cloud-based distributed data center management automation software. Such software is capable of seamlessly migrating IT equipment loads across data centers without affecting the data center SLAs or grid reliability. Petascale and Exascale computing systems will be capable of responding to automated DR signals from electric utilities or ISOs and dynamically shift processing and data storage loads across geographically distant grids. Such capabilities will enable tighter integration of supply-side resources with data centers and their non-fossil-fuel-based local generation sources. The authors recognize that the data set in this study is fairly small. A detailed case-by-case study is recommended to determine the sequence of operations and restoration for the data centers participating in the DR events.

Acknowledgements

This work was sponsored by the Lawrence Berkeley National Laboratory's Demand Response Research Center (http://drrc.lbl.gov/), which is funded by the California Energy Commission's Public Interest Energy Research Program, under Work for Others Contract No. 500-03-026, Am #1 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH1123.

References

[1] Brill, Kenneth G. 2008. Special Report: Energy Efficiency Strategies Survey Results. Uptime Institute.

[2] Koomey, Jonathan 2011. Growth in data center electricity use 2005 to 2010.

[3] EPA. 2007. EPA Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431. Available at http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency and http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf. Both accessed on April 26, 2012.

[4] Flex your Power. 2008. Demand Response Programs. Available from http://www.fypower.org/flexalert/demand_resp_faq.html.

[5] Ghatikar, G., M. A. Piette, S. Fujita, A. T. McKane, J. Q. Han, A. Radspieler, K. C. Mares, and D. Shroyer. 2010. Demand Response and Open Automated Demand Response Opportunities for Data Centers. LBNL-3047E.

[6] Ghatikar, G., V. Ganti, N. Matson, and M. A. Piette. 2012. Demand Response Opportunities and Enabling Technologies for Data Centers: Findings from Field Studies. LBNL-5763E.pdf

[7] Lekov, A., L. Thompson, A. McKane, A. Rockoff, and M. A. Piette. 2009. Opportunities for Energy Efficiency and Automated Demand Response in Industrial Refrigerated Warehouses in California. LBNL-1991E.

[8] McKane A. T., M. A. Piette, D. Faulkner, G. Ghatikar, A. Radspieler Jr., B. Adesola, S. Murtishaw, and S. Kiliccote. 2008. Opportunities, Barriers, and Actions for Industrial Demand Response in California. LBNL-1335E.

[9] Motegi, N., M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu. 2007. Introduction to Commercial Building Control Strategies and Techniques for Demand Response. California Energy Commission. LBNL-59975.

[10] NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 1.0. January 2010. http://www.nist.gov/public_affairs/releases/upload/smartgrid_interoperability_final.pdf

[11] NIST Framework 2.0: NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 2.0, NIST special publication 1108R2. February 2012.

[12] NIST Smart Grid Testing and Certification Committee (SGTCC). 2011. "Interoperability Process Reference Manual, Version 1.0"

[13] GridWise Architecture Council (GWAC), Interoperability Framework, version 0.5. January 2007. http://www.gridwiseac.org/pdfs/interopframework_v05%20070129.pdf

Author Biographies

Venkata Vish Ganti is a Senior Research Associate at LBNL. His research areas include Data Centers, Residential and Commercial Demand Response, and their links to Smart Grid. Mr. Ganti has also been carrying out research on demand response in international markets such as India.

Girish Ghatikar is a Program Manager with at LBNL and leads the data center DR program. He is active in DR and Open Auto-DR (OpenADR) automation and standards activities within U.S. and international Smart Grid activities. Mr. Ghatikar is also the Vice Chairman of the OpenADR Alliance.